EXHIBIT 2

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

CORRIGENT CORPORATION, Plaintiff, v. DELL TECHNOLOGIES INC. and DELL INC., Defendants.	C.A. No. 22-496 (RGA) **JURY TRIAL DEMANDED**
CORRIGENT CORPORATION, Plaintiff, v. ARISTA NETWORKS, INC., Defendant.	C.A. No. 22-497 (RGA) **JURY TRIAL DEMANDED**

DECLARATION OF DR. JAMES OLIVIER
IN SUPPORT OF CORRIGENT'S OPENING CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

TABLE OF EXHIBITS[1]

Exhibit No.	Description
2A	Curriculum Vitae
2B	U.S. Patent No. 5,420,985
2C	RFC 4026, "Provider Provisioned Virtual Private Network (VPN) Terminology" March 2005

[1] Exhibits 2A-2C of my declaration are attached to the Declaration of Andrew Mayo, which is being served on Defendants simultaneously with this document.

I, Dr. James Olivier, declare:

I. INTRODUCTION

1. I have been retained by counsel for Corrigent Corporation ("Corrigent") to evaluate U.S. Patent Nos. 6,957,369 ("'369 Patent") and 7,593,400 ("'400 Patent") (collectively, "patents-in-suit").

2. This Declaration provides opinions relevant to claim construction of certain claim terms or phrases of the '369 and '400 Patents, from the perspective of a person of ordinary skill in the art ("POSA").

II. BACKGROUND AND QUALIFICATIONS

3. I possess the knowledge, skills, experience, training, and education to form an expert opinion and testimony in this case. I am an expert in the field of networking technologies and have been an expert in this field since before 2001. I have designed and developed access systems such as those described in the '369 Patent while working at Marconi Communications and teach MAC bridging and Virtual Private LAN Services like thoses systems described in the '400 patent in my Multiprotocol Label Switching class.

4. I have a Ph.D. from The Ohio State University in Electrical Engineering with minors in Discrete Mathematics, Computer Science, and Microelectronics. My Ph.D. dissertation was based on coding theory and was titled "Concurrent Error Detection in Arithmetic Processors using GAN Codes," in which I developed new codes for use in Arithmetic Processors such as microprocessors. I have published papers on coding theory and continued my work in coding theory for arithmetic processors while at the General Motor Research Laboratory.

5. I have extensive experience in network systems design and development and have specialized in cellular communications product development. I have developed and designed equipment for networks since my time at AT&T Bell Laboratories in the 1990s, where I worked

on AT&T 5ESS, Network Control Point, and Autoplex Series base stations as a Member of the Technical Staff. These systems made extensive use of a counter-rotating ring network topology known as the Common Network Interface. It was there where I first began my work with telecommunication standards bodies as a contributing member of the Asynchronous Transfer Mode ("ATM") ATM Forum. Later at DSC Communications, I was the Senior Manager of the ATM systems engineering group developing ATM packet switches for a new generation base station for Motorola's use in their Centralized Base Station Controller. While at DSC, I was also their corporate representative to the ATM Forum, participating in bi-monthly standards body development. I was one of the first contributors to the wireless standards at the ATM Forum.

6. At Samsung Telecommunications America, I was a Principal Engineer for wireless broadband services over Universal Mobile Telecommunications System ("UMTS"). UMTS is a third-generation ("3G") broadband standard developed by the 3rd Generation Partnership Project ("3GPP"). At Samsung, I worked on designing their next-generation cellular switch, a UMTS mobile switching center ("MSC"). While at Samsung, I was also their corporate representative to the International Telecommunication Union ("ITU"), the United Nations agency responsible for standardizing information and communication technologies. The ITU, under its International Mobile Telecommunications-2000 effort, was part of the 3GPP standards body. It was there that I participated in the development of standards for advanced wireless networks. While at Samsung, I was responsible for the 'services portion" of the UMTS switch. This subsystem provided Internet service, email, and the like to the subscribers.

7. At Marconi, I worked on several systems for the access market, such as Digital Subscriber Line ("DSL") modems and Digital Subscriber Line Access Multiplexers ("DSLAMs"), along with the design of point-to-point wireless systems. At Navini Networks, I was responsible

for layer 2 and layer 3 network protocols for their Wideband Code Division Multiplexed Access ("WCDMA") wireless base stations and broadband modems. These layers were responsible for packet transmissions for various services over the WCDMA air interface. I was also responsible for the development of a Single Sign On ("SSO") system which allowed users to gain information relating to their wireless subscription and delete or add services.

8. I am an Adjunct Professor in the Network Engineering Graduate Program at Southern Methodist University's School of Electrical Engineering, where I teach MPLS networked-enabled applications. Among the MPLS applications I teach is the use of MPLS in advanced network design, including content delivery networks, (CDNs) service provider network design and Intranet vs Internet network designs. I am also the owner of Olivier Consulting, where I provide consulting services for advanced network/product design along with Intellectual Property consulting.

9. I am also currently a Research Professor at the Hunt Institute for Engineering and Humanity, which is part of the Lyle School of Engineering at Southern Methodist University in Dallas, Texas. Here, I also serve as the Program Lead for Transformational Technology. Among the transformational technologies I investigate at the Hunt Institute is the Internet of Things ("IoT"), which makes extensive use of Fourth-Generation and Fifth-Generation networks.

10. I am a co-inventor of U.S. patent No. 8,334,775 issued Dec. 18, 2012 and entitled "RFID-Based Asset Security and Tracking System, Apparatus and Method." This invention relates to a Radio-Frequency Identification (RFID)-based Global Positioning System (GPS) tracking system to provide control and security of assets. This system integrates an RFID-based detection system with a conventional GPS tracking system. The GPS tracking system includes GPS receivers connected over a network to a centrally based GPS monitoring system.

11. A copy of my curriculum vitae with my publications and technical consulting experience is provided as **Ex. 2A**.

III. COMPENSATION

12. I am being compensated for my time at my standard consulting rate, which is $650.00 per hour. My compensation is not dependent in any way upon the outcome of this matter and in no way affects the substance of my Declaration.

IV. MATERIALS CONSIDERED

13. In preparing this Declaration, I have reviewed the '369 and '400 Patents and their respective prosecution histories, Corrigent's Complaints against Defendants Dell Technologies Inc. and Dell Inc. (collectively, "Dell") and Arista Networks, Inc. ("Arista") (including all exhibits attached thereto), the Joint Claim Construction Chart (including the intrinsic evidence cited therein), and all the documents mentioned in this Declaration. Additionally, I am aware of information generally available to, and relied upon by, persons of ordinary skill in the art at the relevant times, including technical dictionaries and technical reference materials (including textbooks, manuals, and technical papers and articles). I also relied on my education, knowledge, and experience and considered the level of ordinary skill in the art as discussed below.

V. PERSON OF ORDINARY SKILL IN THE ART

14. Counsel for Corrigent has informed me that a person of ordinary skill in the art is a hypothetical person who is presumed to have known the relevant art at the time of the invention. In my opinion, the relevant field of art for the patents-in-suit is that of networking configuration, protocols, and switching.

15. I have been informed by counsel for Corrigent that the person of ordinary skill is generally familiar with the type of problems encountered in the field and the prior art solutions to those problems and possesses an ordinary level of creativity.

16. In my opinion, a person of ordinary skill in the art of the patents-in-suit would include someone who had a bachelors or graduate degree in computer science, computer engineering, electrical engineering, or an equivalent training regarding network switching, configuration, and protocols, and approximately two or three years of experience in the field relating to network configuration, protocols, and switching. Lack of work experience can be remedied by additional education, and vice versa.

VI. LEGAL STANDARDS

17. I understand that the words of a claim are generally given their ordinary and customary meaning, which I understand is the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention. I understand that the person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the specification.

18. I understand that claim construction focuses on the "intrinsic evidence," which consists of the claims themselves, the specification, and the prosecution history. I understand that the claims can provide helpful context for how the claim term is used. I understand that the specification is highly relevant to the claim construction analysis and usually dispositive concerning the meaning of a claim term. I further understand, however, that it is inappropriate to import limitations from preferred embodiments recited in the specification into the claim language.

I also understand that, unless the steps of a method actually recite an order, the steps are not ordinarily construed to require an order.

19. I understand that "extrinsic evidence" may also be considered when determining the meaning of a claim. I understand that there are different sources of extrinsic evidence, including dictionaries, inventor testimony, expert testimony, and learned treatises. I understand that intrinsic evidence is generally favored over extrinsic evidence, and that extrinsic evidence may not be used to contradict the meaning of the claim term when read in light of the intrinsic evidence.

VII. '369 PATENT

A. "main module"

20. "Main module" is a term that has a plain and ordinary meaning in the art. It is my opinion that the plain and ordinary meaning of "main module" is a module that has a connection to one or more subsidiary modules via one or more data lines or traces. My opinion in this respect is informed by a review of the intrinsic record and my understanding as a POSA.

21. My understanding, along with the understanding of the POSA, is confirmed by the specification and claim language. The specification refers to "a **main module and at least first and second subsidiary modules, each such module connected to the main module by one or more lines for carrying data**." '369 Patent at 2:42–47 emphasis added; *see also id.* at at Abstract ("subsidiary modules, which are connected to the main module"); 1:27-29 ("a backplane may be used to connect a main module, having a trunk link to a core network, to a number of subsidiary modules, such as line cards"); 1:61-64 ("an electronic system comprises a master module and multiple subsidiary modules, each of which is connected to the master module by one or more connection lines") 3:57–58 ("the main module including a switch having ports for connection to the traces of the backplane;"), 4:60–61 ("[m]ain module comprises a switch and multiple ports.").

Likewise, the claim language states that "the main module include[es] a switch having ports connected to the lines . . . connecting to ports of subsidiary modules via the backplane traces" (claim 1 at 7:6–7) and that "the main module comprises a switch, which is coupled to connect the lines to a network communication trunk, and wherein configuring the main module comprises configuring the switch to connect at least one of the lines of the first subsidiary module to at least one of the lines of the second subsidiary module." (claim 11 at 8:23–24). This underscores that "subsidiary modules" are those that are interconnected with "main modules" in the context of the accused networking systems.

22. Dell and Arista's proposed construction of "a hardware module containing a switch for multiplexing among the subsidiary modules" is too narrow for numerous reasons, and the POSA would not understand the plain and ordinary meaning of "main module" to be so narrow. First, the requirement that a main module "contain[]" a switch is far too narrow. The claim language confirms as much. The independent claims are directed toward "electronic systems" that contain "main modules" and "subsidiary modules." For example, independent claim 8 is directed toward such a system, and does not mandate that a "main module" include a switch. Yet, claim 11, which depends from claim 8, recites that "the main module comprises a switch." This demonstrates that the main module ***can comprise***, but need not comprise, a switch in the context of the asserted patents. The same contrast is apparent from reviewing claims 21 and 23—dependent claim 23 recites that the main module "comprises a switch," while independent claim 21 does not. Further, the specification expressly states that the requirement is optional, as it states that "[t]he main module ***typically*** includes a switch for multiplexing among the subsidiary modules." '369 Patent at 1:30-32 emphasis added.

B. "subsidiary module"

23. "Subsidiary module" is a term that has a plain and ordinary meaning in the art. It is my opinion that the plain and ordinary meaning of "subsidiary module" is a module that has a connection to a main module via one or more data lines or traces, consistent with my opinion for "main module" above. My opinion in this respect is informed by a review of the intrinsic record (including the materials I cited for the "main module" claim element) and my understanding as a POSA.

24. Dell and Arista's proposed construction of "a hardware module under the control of the main module" is too narrow. The POSA would understand that while a main module can control a "subsidiary module," that claim language does not require as much. For example, claim 15 recites a "system control processor," and that system control processor is recited as a distinct claim element from the "main module." Claim 16, which depends from claim 15, states that the first subsidiary module is configured "under control of the system control processor." This suggests that to the extent that the patentee contemplates ***anything*** controlling subsidiary modules, it is the system control processor—not necessarily a "main module." The POSA would understand that only certain embodiments of the invention would have subsidiary modules that are under the control of a main module.

C. "Backplane"

25. "Backplane" is a term that has a plain and ordinary meaning in the art. In general, it is my opinion that the plain and ordinary meaning of "backplane" is "hardware used to establish interconnections between modules." My opinion in this respect is informed by a review of the intrinsic record and my understanding as a POSA.

26. Dell and Arista's proposed construction of "a printed circuit board located in the back of communications equipment containing slot connectors into which edge connectors of

circuit boards are inserted" is too narrow because that is but one example of a backplane implementation. Hardware modules may be interconnected via a backplane in different ways, and not necessarily by plugging printed circuit board ("PCB") cards into a backplane. For example, U.S. Patent No. 5,420,985 to Cantrell, a cited reference on the face of the '369 Patent, describes the use of "Futurebus+," an implementation of backplanes in the prior art that I personally worked on that did not require the use of a PCB or particular hardware for the backplane. *See* Ex. 2B, U.S. Patent No. 5,420,985 to Cantrell at 1:13–22 ("Futurebus+ is an IEEE specification for high-performance backplane-based computing," and "a comprehensive architectural specification . . . for which there are no preconceived restrictions in terms of architecture, microprocessor, and software implementations."). In the Futurebus+ implementation, modules could be interconnected using a backplane by using this backplane in either a loosely coupled or tightly coupled paradigm. *See* Ex. 2B, U.S. Patent No. 5,420,985 at 1:35–37 ("Both loosely coupled and tightly coupled compute paradigms are supported via the parallel protocols and in the message-passing and cache-coherence protocols."). Dell and Arista's proposed construction only allows for tightly coupled backplanes.

27. Dell and Arista's proposed construction is also too narrow because it further attempts to limit backplane's to only being located in the back of communication equipment and only use slot connectors with edge connectors of circuit boards inserted in the slot connectors. Backplanes are not limited to only being located in the back of communication equipment or the use of slot/edge connectors. In the context of the specification and claims, which do not depend on any particular backplane implementation, a backplane can represent any hardware that facilitates interconnections so that communications can be made between components or modules of a switch or networking system. Further, the specification expressly states that the use of slot

connectors into which edge connectors are inserted, e.g. plugged into, is optional for a backplane, as it states that "***Typically, though not necessarily***, the master and subsidiary ***modules plug into a backplane***, …" '369 Patent at 1:64–67 emphasis added.

D. "idle line / idle trace"

28. "Idle line" is a term that has a plain and ordinary meaning in the art. It is my opinion that the plain and ordinary meaning of "idle line" is a line/trace that have spare capacity for testing. My opinion in this respect is informed by a review of the intrinsic record and my understanding as a POSA.

29. My understanding in this respect is confirmed by the context of the specification and claim language. The specification states that idle lines contemplated for testing are those that function "without intruding on normal traffic carried by the system's active lines." '369 Patent at 2:26–29. The point of the testing as being non-disruptive or non-intrusive is stated repeatedly throughout the specification. '369 Patent at 1:5-7, 1:56-60, 4:44-47. The POSA would readily understand that the purpose of testing using "idle lines" are to ensure that data traffic is not disrupted, and the data lines that have spare capacity and can be utilized for testing without disrupting data transmission. The POSA would also understand that lines can be "idle" or not at different points of time, depending on the availability of spare capacity on the lines (i.e., capacity that is not flooded or not at full capacity), as confirmed by the patent specification. '369 Patent at 5:26-30 ("Typically, an operator of system 20 installs subsidiary modules 24 with some spare capacity, to allow provisioning of additional services when requested by users. Therefore, some of traces 34 on backplane 26 may be idle at any given time.").

30. Dell and Arista's proposed construction is also too narrow because it further attempts to limit idle line to a line that is "determined to be inactive." There is nothing in the specification that suggests that the system needs to know, or make any determination, that a

particular line is idle. The point of the invention is simply to test idle lines to facilitate the detection of failures in the background while a switch or router is running.

E. "a system control processor"

31. "[a] system control processor" is a term that has a plain and ordinary meaning in the art. It is my opinion that the plain and ordinary meaning of "[a] system control processor" is exactly what the term means, which is a system control processor. My opinion in this respect is informed by a review of the intrinsic record and my understanding as a POSA.

F. "Order of Steps"

32. I understand that both Dell and Arista wish to require that claims 1, 8, 15, and 21 of the '369 Patent be construed to be limited to particular orderings of steps. I have been asked to render an opinion as to whether these claims require an order of steps. In order to provide my opinion I reproduce claim 1 below:

The '369 Patent
1[p]. In an electronic system that includes a main module and at least first and second subsidiary modules connected to the main module by one or more lines for carrying data, at least some of which lines are sometimes idle, the main module including a switch having ports connected to the lines, a method for self-testing, comprising:
[1A] selecting a first idle line among idle lines connecting the first subsidiary module to a first port of the switch on the main module to serve as an aid line;
[1B] instructing the first subsidiary module to loop back traffic reaching the first subsidiary module via an aid line;
[1C] selecting for testing a second idle line among the idle lines connecting the second subsidiary module to a second port of the switch on the main module;
[1D] configuring the switch to link the first and second ports;
[1E] transmitting test traffic over the second idle line from the second subsidiary module to the main module, wherein the test traffic is conveyed via the switch to the aid line connecting to the first subsidiary module; and
[1F] reporting that a failure has occurred if the test traffic does not return to the second subsidiary module within a predetermined period of time.

33. I have reviewed steps [1A] through [1F] of claim 1 of the '369 Patent. It is my opinion that claim 1 requires the selection of two idle ports (in any order) ([1A] and [1C]), establishing some kind of linkage between the ports [1D], creating a loopback interface [1B], running test traffic between two modules [1E], and reporting failures when traffic does not return to a module in a predetermined time [1F]. The purpose of performing testing could be achieved under numerous orderings and permutations that are not necessarily in the order [1A] through [1F]. For example, it does not matter what order in which the "first" and "second" ports are selected, so steps [1A] and [1C] can occur at different points in time; step [1B] (setting up a loopback) can occur at numerous different points in time, as this step is only needed at the point failures are being checked; and step [1D] (linking ports) can occur at any time before test traffic is run. Therefore, steps [1A] through [1D] can occur in numerous different orders in ways that would allow the purpose of the invention (testing for failures in step [1F]) to be achieved.

34. It is also my opinion that several steps of this claim can occur simultaneously because several of the method steps recited in claim 1 are merely configuration settings that would be set up in connection with a network test. They are steps that could be selected by a user (for example in a configuration or user interface), and implemented by the system simultaneously. More specifically, the POSA would understand that, after inputting settings, the operating system of a network switch could be configured to select ports for testing and link the ports simultaneously (selecting and configuring of steps [1A], [1C], and [1D]). It would also be understood that a port and a loopback interface for a module could be configured to be setup simultaneously (selecting and instructing steps [1A], [1B]). It would even be understood that, once the configuration options are set in the system, steps [1A] through [1D] could all occur simultaneously as the user sets up and executes the testing algorithm, with the running of transmission traffic [1E] and reporting of

failures in step [1F] occurring almost immediately thereafter. Therefore, it is my opinion that Dell and Arista's construction suggesting that step [1B] must necessarily come after [1A]; and that step [1D] must come after steps [1A] and [1C] is incorrect.

35. The other asserted claims do not require the specific order of steps suggested by Defendants for the same reasons set forth above with respect to claim 1.

36. Independent claims 15 and 21 are apparatus claims and do not require specific method "steps" in order to meet the claims. These claims therefore do not require the specific order of steps suggested by Defendants for this additional reason.

VIII. '400 PATENT

A. "forwarding database (FDB)"

37. A forwarding database, also known in the art as MAC Forwarding table or a "forwarding information base" is a common term in computer networking that, in my opinion, the POSA would be familiar with. A forwarding database is a data structure within a router or switch that stores information about how to forward a data packet.

38. My understanding is that Dell does not seek a construction for this term. I understand that Arista contends that the claim term "forwarding database (FDB)" in claims 1 and 11 means "a database shared among the processing channels of a line card holding records associating MAC addresses with the ports of a network node." In my opinion, this construction is too narrow, inserting additional requirements to the claim, such as that the database is "shared among processing channels."

39. When discussing Figure 2, the '400 describes that "[t]he FDB is shared among the processing channels on line card 32," but that is one embodiment from the specification. '400 Patent, 6:66–67. And as I mentioned above, I understand that it is inappropriate to import limitations from embodiments described in the specification into the claim language. I do not see

anything in the claim language that would cause the POSA to limit their understanding of forwarding database to one that is "shared among processing channels" where the claims do not discuss processing channels.

40. Moreover, the claim language already specifically provides that "a respective forwarding database (FDB) *to hold records associating MAC addresses with ports of said plurality of ports of said network node*," (claim 1) and "a respective forwarding database (FDB) *to hold records associating MAC addresses with ports of said plurality of ports of said line cards*," (claim 11) and thus, much of the remainder of Arista's proposed construction is duplicative of the explicit claim language.

41. Finally, I note that Arista's construction introduces "a line card" and "a network node," which is not how I think the POSA would understand the term. The POSA would understand the term within the context of the claim which describes "providing for each of *said member line cards*" and "*said* network node" (claim 1) and a network node "comprising … member line cards … having a respective forwarding database (FDB)" (claim 11). So, the introduction of another line card and another network is not how I believe the POSA would understand the term "forwarding database (FDB)."

B. "said FDB"

42. I understand that Dell and Arista both contend that the claim term "said FDB" in claims 1 and 11 of the '400 Patent is indefinite. I understand that a patent is invalid for indefiniteness if its claims, read in light of the specification delineating the patent, and the prosecution history, fail to inform, with reasonable certainty, those skilled in the art about the scope of the invention.

43. I have reviewed the intrinsic record associated with the '400 Patent with a focus on this claim term. In my opinion, the POSA reading claims 1 and 11 in light of the specification and

intrinsic record would have understood the "said FDB" claim language to have been reasonably certain because it simply refers to the FDB of the first line card, the first line card being the card that is configured to transmit the data packet "to said MAC destination address."

44. Claims 1 and 11 recite "said FDB" three times. In two of the three times, the claims expressly recite "said FDB *of the first line card*." (emphasis added). As such, in view of this context as well as the description provided throughout the specification and claims, I believe that the POSA would understand that the third reference of "said FDB" could also include the FDB "of the first line card."

45. I note that the '400 Patent describes that "[e]ach of the line cards may typically serve as both ingress and egress for data packets and has a respective MAC forwarding database (FDB) that is shared by the ingress and egress functions." Thus, the POSA would understand with reasonable certainty that claims 1 and 11 describe that each LAG member line card includes at least one of its own forwarding databases, and checking the MAC destination address and checking the MAC source address could involve checking of the same FDB, which is on "the first line card."

C. "virtual media access control (MAC) bridge"

46. I understand that Dell and Arista contend that the claim phrase "virtual media access control (MAC) bridge" in claims 8, 11, and 18 of the '400 Patent is indefinite. I again understand that a patent is invalid for indefiniteness if its claims, read in light of the specification delineating the patent, and the prosecution history, fail to inform, with reasonable certainty, those skilled in the art about the scope of the invention.

47. I have reviewed the intrinsic record associated with the '400 Patent with a focus on this claim term. It is my opinion that the POSA would understand the meaning of "virtual media access control (MAC) bridge" with reasonable certainty in the context of the '400 Patent. The term "virtual MAC bridge" is recited in dependent claim 8, where the claim language establishes that

"the network node [can be] configured to operate as multiple virtual MAC bridges in a Layer 2 virtual private network, wherein each virtual MAC bridge is configured to serve a respective VPN instance." '400 Patent at 11:65–12:5 (emphasis added). Substantially identical language appears in the specification, and states that "each virtual MAC bridge is configured to serve a respective VPN instance." *Id.* at 4:54–64 (emphasis added). Claim 11 also recites "so that the node operates as a virtual media access control (MAC) bridge in said Layer 2 data network." *Id.* at 12:30–31. The specification also describes virtual LANs and how "a given VPLS instance may be partitioned into a number of virtual LANs (VLANs), which generally operate in the manner defined in the above-mentioned IEEE Standard 802.1Q." *Id.* at 7:5–8. The POSA would understand, with reasonable certainty, that the construction of "virtual media access control (MAC) bridge" is a media access control (MAC) bridge that serves a virtual private network (VPN) instance.

48. Such virtual switches were well-known and had a well-understood meaning in the art, as reflected, for example, in RFC 4026. *See* Ex. 2C, RFC 4026. RFC 4026, issued March 2005, is titled "Provider Provisioned Virtual Private Network (VPN) Terminology" and underscores that the virtualization of MAC bridges and other components of networks were well known as of the priority date. Specifically, RFC 4026 refers to a Virtual Private Lan Service that "emulates a learning bridge, and forwarding decisions are taken based on MAC addresses." *Id.* at p.4. This reinforces that the meaning of this term, with reasonable certainty, is a "media access control (MAC) bridge that serves a virtual private network (VPN) instance," consistent with the plain and ordinary meaning of this claim term. The POSA would understand that this construction applies not just to claim 8, but to claims 11 and 18 as well.

D. "Order of steps"

49. I understand that both Dell and Arista seek to require that claim 1 of the '400 Patent be construed to be limited to particular orderings of steps. I have been asked to render an opinion

as to whether claim 1 requires an order of steps. In order to address any alleged relation between steps and provide my opinion, I reproduce method claim 1 below:

The '400 Patent
1. A method for communication, comprising:
[1A] configuring a network node having a plurality of ports, and at least first and second line cards with respective first and second ports, to operate as a distributed media access control (MAC) bridge in a Layer 2 data network;
[1B] configuring a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports and a plurality of conjoined member line cards;
[1C] providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;
[1D] receiving a data packet on an ingress port of said network node from a MAC source address, said data packet specifying a MAC destination address on said Layer 2 data network;
[1E] conveying, by transmitting said data packet to said MAC destination address via said first port, said received data packet in said network node to at least said first line card for transmission to said MAC destination address;
[1F] if said MAC destination address does not appear in said FDB, flooding said data packet via one and only one LAG port of said plurality of LAG ports;
[1G] checking said MAC source address of the data packet against records in said FDB of said first line card; and
[1H] if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card of said plurality of member line cards.

50. From the Joint Claim Construction Chart, I understand that Dell and Arista's position is that [1E] (conveying the received data packet to the egress line card) must occur before [1F] (destination address checking/subsequent flooding) and [1G] (source address checking). It is my opinion that there is no language in the claims or specifications that require this ordering.

51. Step [1F] is essentially the MAC destination address checking step. No claim language in step [1F] provides any sequence relative to [1E], [1G], or [1H]. In my opinion, [1F]

happens after [1D], where the packet is received, because until the packet is received, the system cannot check the destination address specified by the data packet. In my opinion, the POSA would understand that [1F] most likely happens *before* [1E], not after as suggested by Dell and Arista. Examples in the specification describe checking the destination address as part of the ingress processing. For example, the '400 Patent provides:

> When an ingress line card receives an incoming data packet over the VPN on one of its ports, it consults the FDB in order to choose the line card and port through which the packet should be forwarded based on the MAC destination address (or floods the packet through the ports in the VPN when the MAC destination address does not appear in the FDB).

'400 Patent, 3:10–16; *see also* 1:38–40; 7:26–30 (when packet processor 52 "extracts the other key parameters ***from the incoming packet (including the MAC destination address*** (DA).").

52. In my opinion, Dell and Arista's proposal would make little sense to the POSA. How would the system know where to "convey" the packet [1E] if the destination address is not checked *before* the conveying? If the packet has already made it to the specific egress card (*i.e.*, the first line card in [1E]) for transmission to the MAC destination address, why would the system then check the FDB for the MAC destination address and then initiate a flooding operation? In my opinion, based on the materials I have reviewed to date, [1F] is not required to happen before [1E], and if anything, the POSA would understand that it is more likely that the checking portion of [1F] happens before [1E].

53. Step [1G] and [1H] together constitute the MAC source address learning and synchronization. Once again, it is my opinion that there is no language in claim 1 that requires the MAC source address learning and synchronization to occur *after* [1E] as Dell and Arista propose. One way the source address learning is different than the destination address that would be apparent to the POSA is that the source learning is important for the *next data packet* destined to

the source address, whereas the destination address checking is an effort to handle and process the current, *received data packet*. In other words, the MAC source address learning and synchronization is an advancement that prevents flooding of future packets—the claimed messaging prevents the "incomplete" FDB problem the patent describes. *See* '400 Patent, 3:34–53.

54. While the '400 Patent describes that "[i]n some embodiments … [t]he egress line card (or line cards) that is to transmit the packet onward checks the MAC source address of the data packet against the records in its own FDB," (3:3–20), there is nothing in the specification that disclaims or narrows the claims in a way that *requires* [1G] to occur after [1E]. Moreover, while there was a technical, logical data flow reason for the POSA to understand that MAC destination address checking of the received data packet would occur before "conveying" (*i.e.*, if the destination address isn't checked, how does the system know which card is the correct egress card), the same is not true for the MAC source address learning in claim 1. Indeed, the general notion of MAC learning would be familiar to a POSA. See, e.g., Background of the '400 Patent, 1:33–38. As such, in my opinion, the POSA would understand that [1G] could happen before [1E].

55. Dell and Arista also provide that step [1F] "must be performed before the record is added to the FDB (1[h])." JCCC at 15. As described above, [1F] involves the checking of the destination address, but [1H] involves adding a new record "if said FDB of said first line card does not contain a record of an association of ***said MAC source address*** with said ingress port." [1H] has nothing to do with the destination address. If there is no record of the *destination address* on the other hand, the method includes flooding the packet. In that scenario, the location of the destination address is unknown—there is no association to add. Once again, it is my opinion that the POSA would find Dell and Arista's proposal illogical.

56. I agree with Dell and Arista that the POSA would understand that [1H] (the source address learning) must happen after [1G], as the plain language of [1H] makes it contingent on the FDB of the first line card not containing a record of the source address/ingress port association.

E. "conveying … said received data packet … to at least said first line card for transmission to said MAC destination address / said ingress line card conveys said data packet … to at least said first line card for transmission to said MAC destination address"

57. This claim phrase is the "conveying step" I have discussed a bit already above. While Dell and Arista propose different constructions, the difference is a slight change in the word ordering, but both of their proposed constructions are similar. However, in my opinion, the plain language of the claims would be readily understood by the POSA without any additional construction.

58. For example, claim 11 plainly states:

> said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address.

59. In my opinion, the POSA would easily understand this language without the need or desire for any construction. Dell and Aritsa's constructions unnecessarily introduce "a line card containing the ingress port" to both claims 1 and 11. But a POSA would not consider new language, but instead would rely on the express language of the claims themselves. Regarding claim 11, the plain language already says the "ingress line card conveys." This would be clear to the POSA. Regarding claim 1, the plain language provides for conveying the "*received* data packet," where the data packet was received on an ingress port (*see* limitation [1D]). I see no benefit provided to the POSA by Dell and Arista's proposed constructions.

60. I also note that Dell and Arista's constructions eliminates claimed structure from claim 11. Claim 11 cites that the node comprises "a switching core" and that the data packet is

conveyed "via said switching core to at least said first line card." Dell and Arista appear to eliminate this language from the claim.

61. In my opinion, the POSA would understand the plain language of the claims without need for any additional construction.

F. "providing for each of said member line cards a respective forwarding database (FDB)"

62. My understanding is that Dell does not seek a construction for this phrase. I understand that Arista contends that the claim phrase "providing for each of said member line cards a respective forwarding database (FDB)" from claim 1 means "providing for each of said member line cards a single forwarding database (FDB)." The only difference between the claim language and Arista's proposal is replacing the word "respective" from the claim with "single." In my opinion, the POSA would not understand "respective" FDB to mean "single" FDB in this context.

63. The POSA would understand that "respective" is describing that *each* member line card includes its own FDB. But the POSA would understand that while each line card might include at least one FDB, the claim is not limited to an instance where each line card is limited to only a single FDB.

64. In my opinion, replacing "respective" with "single" is inconsistent with use of the word "respective" throughout the claims and specification in other contexts:

- "at least first and second line cards with *respective* first and second ports" (claim 1)
- "the *respective* ports" (claim 2)
- "each line card having *respective* ports and having a respective forwarding database (FDB) to hold records associating MAC addresses with said *respective* ports" (claim 11).

In each of these instances, it would not make sense to the POSA to replace "respective" with "single" as Arista has proposed for the FDB aspect of claim 1.

* * *

65. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Date: January 9, 2024

James L. Olivier, Ph.D.